

Fast and Vectorizable Alternative to Binary Search in $O(1)$ Applicable to a Wide Domain of Sorted Arrays of Floating Point Numbers

Fabio Cannizzo¹

Abstract

Given an array X of $N + 1$ strictly ordered floating point numbers¹ and an array Z of M floating point numbers belonging to the interval $[X_0, X_N)$, a common problem in numerical methods is to find the indices of the largest numbers in the array X which are smaller or equal than the numbers in the array Z . The general solution to this problem is to call M times the well known *binary search* algorithm, which has complexity $O(\log_2 N)$ per individual search and, in its classic formulation, is not vectorizable, i.e. it is not SIMD friendly. This paper describes technical improvements to the binary search algorithm, which make it faster and vectorizable. Next it proposes a new vectorizable algorithm applicable to a wide set of X partitions, which is based on an indexing technique and allows to solve the problem with complexity $O(1)$ per individual search at the cost of introducing an initial overhead to compute the index and requiring extra memory for its storage. Test results using streaming SIMD extensions compare the performance of the algorithm versus various benchmarks and demonstrate its effectiveness. Depending on the test case, the algorithm performs up to 43 times faster than the classical binary search in single precision and 39 times faster in double precision. Applicability limitations and cache-friendliness related aspects are also discussed.

Keywords: interpolation, spline, SIMD, SSE, vectorization, lower_bound

1. Introduction

Given an array X of strictly increasing floating point numbers $\{X_i\}_{i=0}^N$ and a set of floating point numbers $\{Z_j\}_{j=1}^M$ belonging to the interval $[X_0, X_N)$, a common problem in numerical methods algorithms is to find the indices of the largest numbers in the array X which are smaller or equal than the numbers Z_j .

This problem arises for instance in the context of piece-wise interpolation, where a domain $[X_0, X_N)$ is partitioned in sub-intervals $\{[X_i, X_{i+1})\}_{i=0}^{N-1}$ and different interpolation functions $g_i(x)$ are associated with each sub-interval. To compute the interpolated value for a number Z_j , the index i of the sub-interval containing it needs to be resolved first.

¹DISCLAIMER: Although Fabio Cannizzo is employed by Standard Chartered at the time this paper is written, this paper has been produced by Fabio Cannizzo in a personal capacity and Standard Chartered is not associated or responsible for its content in any way.

¹a floating point number is a rationale number which can be represented exactly on a computer given a chosen floating point representation

This is a special case of the more general problem of searching in a sorted array for a matching element. The general solution to this problem is to call M times the *binary search* algorithm, which has complexity $O(M \log_2 N)$. An history of the *binary search* algorithm is discussed in [3], where the first published mention in the literature is attributed to J. Mauchly [9], who in 1946 proposed a binary search algorithm to find an exact match in a sorted array of length 2^N . It is not until 1960 that a version of the algorithm that works for any N was published by D. Lehmer [10]. The next step was taken by Bottenbruch [11], who presented a variation of the algorithm that avoid a separate test for equality until the very end, thus speeding up the inner loop.

A well known variation of Bottenbruch's algorithm, specialised to find the index of the largest numbers in the array X , which is smaller or equal than some number $z \in [X_0, X_N)$, is described by Press et al. [4] and is referred to in the sequel as the *classical* version of the algorithm. This requires a control flow branch, which incurs penalties on many CPU architectures and is not vectorizable, i.e. it does not benefit from the vectorial capabilities of modern CPUs. Furthermore the algorithm is not cache-friendly as explained in Khuong [8] and Morin et al. [2].

In some special cases, when either the X_i or the Z_j numbers exhibit particular patterns, more efficient algorithms are available. Examples are when the numbers X_i are equally spaced or when the numbers Z_j are sorted, and the problem can be solved with complexity $O(M)$ and $O(M + N)$ respectively. However no generic alternative exists.

Many variations of the *binary search* algorithm have been developed to improve *binary search*. *Fibonacci search*, proposed by J. Kiefer [12] and D. Ferguson [13], has the same complexity as *binary search*, but improve search time when some region of the array can be accessed faster than others. *Exponential search*, proposed by Bentley et al. [14], for a single search has complexity $O(\log_2 i)$ where i is the sought index, but, due to increased computation cost, it becomes an improvement over *binary search* only if the target value lies near beginning of the array. *Ternary search* solves the problem in $\log_3 N$ iterations, but it is still less efficient than *binary search*, because at every iterations it performs two comparisons instead of one, hence the total number of comparisons is $2 \log_3 N \approx 1.26 \log_2 N > \log_2 N$. Several other variations have been analyzed by Morin et al. [2] with the intent to reduce branch mis-prediction and cache misses.

This paper describes technical improvements to two variations of the *binary search* algorithm available in the literature, which makes them slightly faster and vectorizable. The complexity of the algorithms remains $O(M \log_2 N)$, but performance improves by a proportionality factor α/d , where α is a constant smaller than one due to the performance improvements and d is the number of floating point numbers which can be processed simultaneously², assuming perfect vectorization³.

Next it proposes a new vectorizable algorithm based on a indexing technique, which solves the problem with reduced complexity $O(M)$, at the cost of introducing an initial overhead to compute the index and requiring extra memory for its storage. Although the algorithm has fairly wide applicability, there are particular situations where the layout of the numbers X_i might make its use unfeasible. These limitations are analysed and discussed and a variations of the algorithm which mitigate them at the cost of sacrificing some performance are introduced. A cache-friendly version of the algorithm is also proposed.

² d depends on the chosen set of vectorial instructions and floating point representation (e.g. with SSE instructions in single precision $d = 4$)

³perfect vectorization means that all scalar instructions used in the algorithm have a vectorial equivalent

2. Problem Statement

Let:

- X be an array of $N + 1$ floating point numbers sorted in ascending order, i.e. $X_{i-1} < X_i$, $i = 1 \dots N$. The array X is assumed to be stored in a container supporting access by index in constant time. The numbers in X exhibit no special pattern.
- Z be an unordered array of M floating point numbers contained between the first and the last element of the array X , i.e. $Z_j \in [X_0, X_N)$, $j = 0 \dots M - 1$. Note that the exclusion of the last point X_N does not imply any loss of generality, as X_N could simply be replaced with the next larger floating point number.

for each number Z_j the index i of the largest element in the array X such that $X_i \leq Z_j$ is sought. Such array of indices is denoted as $\{I_j\}_{j=0}^{M-1}$.

It is assumed that M is large and memory is not scarce, therefore it is worth investing upfront in a preliminary analysis of the structure of the array X and in the creation of an auxiliary data structure in order to later achieve superior computational performance when computing the indices $\{I_j\}_{j=0}^{M-1}$.

3. Binary Search

3.1. Classical Binary Search

In absence of any special pattern of the values X_i , the classical solution to this problem is the *binary search* algorithm, which in scalar version has complexity $O(\log_2 N)$ per individual search. A simple scalar pseudo-code implementation is given in algorithm 1.

Algorithm 1 Classical Binary Search (scalar implementation)

```

function BINARYSEARCH(input:  $z, \{X_i\}_0^N$ , output:  $i$ )
   $low \leftarrow 0$ 
   $high \leftarrow N$ 
  while  $high - low > 1$  do                                 $\triangleright$  terminating condition depends indirectly on  $z$ 
     $mid \leftarrow (low + high)/2$ 
    if  $z < X_{mid}$  then                                      $\triangleright$  code branch
       $high \leftarrow mid$ 
    else
       $low \leftarrow mid$ 
    end if
  end while
   $i \leftarrow low$ 
end function

```

A close analysis of this implementation highlights the following weaknesses: the body of the loop contains a branch instruction and for an arbitrary z the chances for the boolean condition to be *true* are about 50%, which makes branch prediction algorithms used in modern CPUs ineffective; the algorithm is not easily vectorizable because, for different z the boolean condition may

resolve differently causing the program flow to take different code paths and possibly requiring a different number of iterations for the loop to complete.

Because proper vectorization is not possible, a vectorial implementation is obtained trivially iterating on all elements of the array Z in steps of one element, with complexity $O(M \log_2 N)$.

3.2. Binary Search Revisited

An improved variation of binary search proposed in [1] is based on the observations that some unnecessary extra iterations in the loop do not change the result and that the maximum possible number of potentially required iterations is $\lceil \log_2 N \rceil$.

Let p be the number of bits necessary to represent the number N , which is $p = 1 + \lfloor \log_2 N \rfloor$, b_k be the binary value taken by the k -th bit of the sought index i , $c_k = 2^{k-1}$ and $a_k = b_k c_k$, the sought index i has binary representation $\sum_{k=1}^p a_k$.

The bits of the index can be resolved one by one starting from the highest order one as follows: first, if $z \geq X_{c_p}$ then the p -th bit of the sought index i must be set, i.e. $b_p = 1$, next, if $z \geq X_{a_p+c_{p-1}}$ then the $(p-1)$ -th bit of the index i must be set, i.e. $b_{p-1} = 1$, and so on, the values of the remaining bits are obtained iterating the procedure.

This yields to algorithm 2, where the input argument P is the precomputed constant $P = c_b = \lfloor \log_2 N \rfloor$. The number of iterations is fixed to $p-1$, so no longer dependent on z , however, as the algorithm proceeds it can happen that the candidate index of the vector X exceeds the size of the vector, therefore, it requires a double boolean condition with short boolean evaluation, which makes it difficult to vectorize.

Algorithm 2 Binary Search Revisited (scalar implementation)

```

function BITSETBINARYSEARCH1(input:  $z, \{X_i\}_{i=0}^N, P$ , output:  $i$ ) ▷  $P = 2^{\lfloor \log_2 N \rfloor}$ 
   $i \leftarrow 0$ 
   $k \leftarrow P$ 
  repeat
     $r \leftarrow i \mid k$  ▷ bitwise OR
    if  $r < N \ \&\& \ z \geq X_r$  then ▷ short boolean evaluation
       $i \leftarrow r$ 
    end if
     $k \leftarrow k/2$  ▷ bitwise right shift
  until  $k = 0$ 
end function

```

3.3. Vectorizable Binary Search Revisited

The first boolean condition in algorithm 2 can be avoided with a simple trick. Noting that $Z_j < X_N$ for any j , if the array X is extended to the right side by padding it with the last entry X_N up to a size equal to the largest possible index representable with p bits (i.e. $Q = 2^p - 1$), the condition $z \geq X_r$ would resolve to *false* for any $r > N$ generated by the algorithm. This introduces an extra memory requirement of about $(2^p - N)S$ bytes, where S is the size in bytes of the elements of array X (4 for single precision and 8 for double precision).

This yields algorithm 3, which has still complexity $O(\log_2 N)$, but does not contain any code branch (only a conditional assignment, which can be implemented efficiently with modern CPUs instruction set) and is easily vectorizable.

Algorithm 3 Vectorizable Binary Search Revisited (scalar implementation)

```
function BITSETBINARYSEARCH2(input:  $z, \{X_i\}_{i=0}^Q, P$ , output:  $i$ ) ▷  $P = 2^{\lfloor \log_2 N \rfloor}$   
   $i \leftarrow 0$   
   $k \leftarrow P$   
  repeat  
     $r \leftarrow i \mid k$  ▷ bitwise OR  
    if  $z \geq X_r$  then  
       $i \leftarrow r$  ▷ conditional assignment  
    end if  
     $k \leftarrow k/2$  ▷ bitwise right shift  
  until  $k = 0$   
end function
```

A vectorial implementation is obtained vectorizing the function and iterating on all elements of the array Z in steps of d elements, where d is the number of floating point numbers which can be processed simultaneously⁴. Note that, working with streaming SIMD extensions perfect vectorization is not achievable, because the indices contained in vector r are not contiguous, therefore the d elements of the vector X_r must be fetched from memory sequentially.

If memory is scarce, an alternative to extending and padding the array X is to take at every iteration the maximum between r and N as the index of the element in X used for comparison. This can be resolved by the compiler without branching as a conditional assignment and leads to algorithm 4. This however leads to a loss of performance, which is still superior to the classic *binary search*, but inferior to algorithm 5, which is described in the next section and does not require any extra memory, therefore test results are not reported in section 5.3.

Algorithm 4 Vectorizable Binary Search Revisited With No Padding (scalar implementation)

```
function BITSETBINARYSEARCH3(input:  $z, \{X_i\}_{i=0}^N, P$ , output:  $i$ ) ▷  $P = 2^{\lfloor \log_2 N \rfloor}$   
   $i \leftarrow 0$   
   $k \leftarrow P$   
  repeat  
     $r \leftarrow i \mid k$  ▷ bitwise OR  
     $w \leftarrow \min(r, N)$  ▷ conditional assignment  
    if  $z \geq X_w$  then  
       $i \leftarrow r$  ▷ conditional assignment  
    end if  
     $k \leftarrow k/2$  ▷ bitwise right shift  
  until  $k = 0$   
end function
```

3.4. Offset-Based Binary Search

A quite effective variation of the classic binary search algorithm is proposed in [2], where, as the search range is progressively bisected, its current coordinates are stored as a pair of *start index*

⁴ d depends on the family of streaming SIMD extensions used and on the chosen number representation, e.g. *single* or *double* precision.

and *range size*, as opposite to a pair of *start index* and *end index*. This allows to perform only one conditional assignment to update the *start index*, as the size of the search range decreases deterministically by a factor 1/2. Furthermore, the algorithm is already optimized to be branch-free.

The algorithm presented here contains a technical improvement with respect to the original one proposed in [2], as the number of iterations is fixed and not dependent on the input data (as in algorithm 3), thus allowing for vectorization and helping the branch predictor. Furthermore the problem solved in this paper is slightly different from the one solved in [2], which search for the index i of the element of the array X matching exactly the number z , which leads to more differences.

The pseudo-code is presented in algorithm 5, where the search function receives in input the following pre-computed constants

$$\begin{array}{ll} M = (N + 1)/2 & \text{mid index} \\ S = N + 1 - M & \text{mid size} \\ J = \lfloor \log_2(N + 1) \rfloor & \text{number of iterations} \end{array}$$

Algorithm 5 Offset Based Binary Search (scalar implementation)

```

function BINARYSEARCHOFFSET(input:  $z, \{X_i\}_{i=0}^N, F, S, J$  output:  $i$ )
   $i \leftarrow 0$ 
  if  $z \geq X_F$  then                                      $\triangleright$  Assume at least one iteration, i.e.  $J > 0$ 
     $i \leftarrow F$                                           $\triangleright$  Conditional assignment
  end if
  while  $J > 0$  do
     $J \leftarrow J - 1$ 
     $H \leftarrow S/2$ 
     $F \leftarrow i + H$ 
    if  $z \geq X_F$  then                                      $\triangleright$  Conditional assignment
       $i \leftarrow F$ 
    end if
     $S \leftarrow F - H$ 
  end while
end function

```

A vectorial implementation is obtained vectorizing the function and iterating on all elements of the array Z in steps of d elements, where d is the number of floating point numbers which can be processed simultaneously. Note that, working with streaming SIMD extensions perfect vectorization is not achievable, because the indices contained in vector F of algorithm 5 are not contiguous, therefore the d elements of the vector X_F must be fetched from memory sequentially.

3.5. Cache Friendly Binary Search

Modern CPUs exploit a small amount of integrated high performance memory called cache, which can be accessed much faster than the regular RAM. In simple terms this can be thought of as a temporary copy of various fragments of the RAM. Any data needed by the program is first sought in cache, and, if not found⁵, is uploaded from RAM to cache, then used. As the cache is

⁵when this is happens, it is said that a *cache-miss* event occurs, which has performance penalties

small, every time something is loaded into it, in order to make space some other piece of data must be displaced.

Because cache memory is extremely fast, it is beneficial for a program to carry out as much work as possible on a data set small enough to fit in cache, so that the accesses to RAM, which are expensive, are minimized. An algorithm designed with this property in mind is said to be *cache friendly*.

The degree of efficiency in the use of cache memory for various formulation of comparison based search algorithms has been studied systematically in [2]. The conclusion is that one of the most cache friendly implementation of the search algorithms relies on reordering the array X according with a special layout, called *Eytzinger layout* [15].

The algorithm incurs some setup cost to reorder the array X and requires extra storage space to store it. An efficient implementation requires that the size of the original array X is $2^L - 1$ for some value L . Algorithm 6 describes the index search procedure, assuming that the Eytzinger layout has already been computed and stored in a new array Y . For an explanation of how the Eytzinger layout is organized, the reader is referred to [2].

The problem solved in this paper is slightly different from the one solved in [2], which search for the index i of the element of the array X matching exactly the number z , therefore the algorithm proposed here is also slightly different. Furthermore, the array Y is padded with the last element of the array X_N until it reaches the size $2^L - 1$, so that the algorithm does not require the introduction of checks on the index range and can work at its best efficiency. This trick, already used in algorithm 3, does not affect results, because as part of the problem statement $z < X_N$.

This introduces an extra memory requirement of about $(2^L - N)S$ bytes, where S is the size in bytes of the elements of array X (4 for single precision and 8 for double precision).

The detailed search procedure is listed in algorithm 6, which uses the following precomputed constants:

$$L = 1 + \lfloor \log_2(2 + N) \rfloor$$

$$M = \text{not}(2L) \quad (\text{not is the bitwise not operator})$$

Algorithm 6 Eytzinger Search (scalar implementation)

```

function EYTZINGERSEARCH(input:  $z, \{Y_i\}_{i=0}^{2^{L+1}-2}, M, L$ , output:  $i$ )
   $P \leftarrow 1$ 
  if  $z \geq Y_0$  then
     $P \leftarrow 2$                                  $\triangleright$  conditional assignment
  end if
  while  $L > 1$  do
     $Q \leftarrow 1$ 
    if  $z \geq Y_P$  then
       $Q \leftarrow 2$                                  $\triangleright$  conditional assignment
    end if
     $P \leftarrow 2P + Q$ 
     $L \leftarrow L - 1$ 
  end while
   $i \leftarrow P \& M$                                  $\triangleright$  bitwise AND
end function

```

A vectorial implementation is obtained vectorizing the function and iterating on all elements of the array Z in steps of d elements, where d is the number of floating point numbers which can be processed simultaneously. Note that, working with streaming SIMD extensions perfect vectorization is not achievable, because the indices contained in vector P of algorithm 6 are not contiguous, therefore the d elements of the vector X_P must be fetched from memory sequentially.

4. Direct Search

A scalar algorithm with complexity $O(1)$ per individual search can be obtained via construction of an auxiliary function which maps real numbers $z \in [X_0, X_N]$ directly to the sought indices i according with the following simple procedure.

Let f be a function which maps floating point numbers $z \in [X_0, X_N]$ to natural numbers in $[0, R]$ for some value R and satisfies the following properties

$$f(X_0) = 0 \tag{1a}$$

$$f(X_N) = R \tag{1b}$$

$$\forall a, b \in [X_0, X_N], a > b \implies f(a) \geq f(b) \tag{1c}$$

$$f(X_{i+1}) > f(X_i), \quad i = 0 \dots N - 1 \tag{1d}$$

Let $\{K_j\}_{j=0}^R$ be a sorted array of natural numbers mapping the indices j generated by function f to indices of the array X as follows

$$K_j = \begin{cases} 0, & j = 0 \\ i, & f(X_{i-1}) < j \leq f(X_i) \end{cases} \tag{2}$$

Possible pseudo-code to construct the array K as specified in (2) is proposed in algorithm (7). The definition (2) of array K , implies the following property

$$c = f(X_i) \implies K_c = i \tag{3}$$

Given a floating point number $z \in [X_i, X_{i+1})$, property (1c) of function f guarantees that

$$a = f(X_i) \leq j = f(z) \leq f(X_{i+1}) = b,$$

and property (1d) implies that $a < b$. Because of (3) $K_a = i$ and $K_b = i + 1$, hence the index $t = K_j$ can be either i or $i + 1$

$$i = K_a \leq t = K_j \leq K_b = i + 1$$

Since $z \in [X_i, X_{i+1})$, if $t = i$ then $z \geq X_t$, while if $t = i + 1$ then $z < X_t$, therefore the sought index can be trivially resolved comparing z with X_t :

$$i = \begin{cases} t - 1 & \text{if } z < X_t, \\ t & \text{otherwise} \end{cases}$$

Algorithm 7 Initialization of array K (pseudo-code)

```
function INITK(input:  $\{X_i\}_{i=0}^N, f(z)$ , output:  $\{K_j\}_{j=0}^R$ )  
   $b \leftarrow R$   
   $i \leftarrow N$   
  while  $b \geq 0$  do  
     $t \leftarrow f(X_i)$   
    while  $b > t$  do                                      $\triangleright$  always false at the first iteration, when  $i = N$   
       $K_b \leftarrow j$   
       $b \leftarrow b - 1$   
    end while  
    if  $b = t$  then                                        $\triangleright$  always true at the first iteration, when  $i = N$ , thus  $j$  gets initialized  
       $j \leftarrow i$   
       $K_b \leftarrow j$   
       $b \leftarrow b - 1$   
    end if  
     $i \leftarrow i - 1$   
  end while  
end function
```

4.1. Proposed Choice for the Function f

The search procedure described in section 4 relies on the existence of a function f satisfying properties (1). Since the largest possible value generated by the function is R , which defines the size of the array $\{K_j\}_{j=0}^R$, in order to minimize storage space requirements and initialization cost, it is desirable for R to be as small as possible. Furthermore, since function f is used in the search routine, it is desirable for it to be computationally fast.

A possible choice, not necessarily optimal, is to use the simple formula

$$f(z) = \lfloor H(z - X_0) \rfloor \quad (4)$$

where H is an appropriately chosen constant. This is clearly a computationally efficient function, as its evaluation requires only a multiplication, a subtraction and a truncation.

Condition (1d) requires that H satisfies the inequalities

$$\lfloor H(X_{i+1} - X_0) \rfloor > \lfloor H(X_i - X_0) \rfloor, \quad i = 0 \dots N - 1 \quad (5)$$

and the truncation operation can be removed writing the more restrictive system of inequalities

$$H(X_{i+1} - X_0) > 1 + H(X_i - X_0), \quad i = 0 \dots N - 1 \quad (6)$$

which yields the theoretical lower bound

$$H > \frac{1}{\min_i \{X_{i+1} - X_i\}_{i=0}^{N-1}} \quad (7)$$

Once a value for H satisfying the lower bound (7) is chosen, R is simply determined applying formula (4) to X_N

$$R = \lfloor H(X_N - X_0) \rfloor \quad (8)$$

A possible value for the constant H which strictly satisfy inequality (7) could be obtained as

$$R = 1 + \left\lceil \frac{X_N - X_0}{\min_i \{X_i - X_{i-1}\}_{i=1}^N} \right\rceil \quad (9a)$$

$$H = \frac{R}{X_N - X_0} \quad (9b)$$

In reality, as explained in section 4.4, because of rounding errors this is not a robust approach to determine H and a different methodology must be used.

Note that, because of the transformations performed of inequalities (5) into more restrictive conditions, the lower bound H proposed in (7) is not guaranteed to be optimal, i.e. it is not guaranteed to be the smallest possible value which would make function (4) compatible with properties (1). For instance, given the array $X = \{0, 0.5, 0.7, 1.1\}$, equations (7) yields $H \approx 6.36$, however smaller values would also be admissible, e.g. $H = 3$. The determination of the optimal value of H satisfying properties (1) is not addressed in this paper.

4.2. Algorithm

Summarizing, using the choice for function f proposed in (4), for a given number $z \in [X_0, X_N)$ the index i such that $z \in [X_i, X_{i+1})$ can be obtained with the following procedure, given in pseudo-code in algorithm 8:

1. compute the index j using (4)
2. read the correspondent index i stored in K_j
3. verify if the numbers z is to the right or to the left of X_i , and, if it is to the left, decrement the index i by one.

Note that the last step does not involve any conditional jump. In the scalar case it can be resolved via conditional assignment. In the vectorial case with streaming SIMD extensions a floating point comparison operations returns a bit mask, which, if reinterpreted as a signed integer, is either 0 or -1 and can be trivially added to i .

Algorithm 8 Direct Search (scalar implementation)

```

function DIRECTSEARCH(input:  $z, \{X_i\}_{i=0}^N, \{K_j\}_{j=0}^R, H$  output:  $i$ )
     $j \leftarrow \lfloor H(z - X_0) \rfloor$ 
     $i \leftarrow K_j$ 
    if  $z < X_i$  then
         $i \leftarrow i - 1$  ▷ conditional assignment
    end if
end function

```

Similarly to what already discussed for algorithm 3, working with Intel SIMD instructions perfect vectorization is not achievable, as fetching from memory the relevant elements of K_j and X_i must be done sequentially.

4.3. Geometric Interpretation

The algorithm described in section (4.2) has an intuitive equivalent geometric interpretation. Let $\{Y_j = X_0 + j/H\}_{j=0}^{R+1}$ be a fictitious array of equally spaced real numbers, the segments $\{[Y_j, Y_{j+1})\}_{j=0}^R$ have constant length $1/H$ and overlap with the segments $\{[X_i, X_{i+1})\}_{i=0}^{N-1}$.

The function $j = f(z)$ defined in (4) can be interpreted as the mapping from a real number $z \in [X_0, X_N)$ to the interval $[Y_j, Y_{j+1})$ which contains it.

The array $\{K_j\}_{j=0}^R$ defined in (2) can be interpreted as the mapping of the abstract segments $\{[Y_j, Y_{j+1})\}_{j=0}^R$ to elements of the array X as described below and illustrated in figure 1

$$K_j = \max_i \{i : X_i \leq Y_{j+1}\}^N$$

Given a real number $z \in [X_0, X_N)$, the operations $j = f(z)$ and $t = K_j$, identify respectively the segment $[Y_j, Y_{j+1})$ containing z and an associated element X_t . Because of (7), it is guaranteed that the segment $[Y_j, Y_{j+1})$ is smaller than the smallest segment $\min_i \{[X_i, X_{i+1})\}_{i=0}^{N-1}$, therefore it can contain at most one single element of the array X . This means that either $Y_j \leq X_t < Y_{j+1}$, implying that z could be either to the left or to the right of X_t , or $X_t < Y_j$, implying that z is always to the right of X_t . In both cases, the sought index i is either t , if $z < X_t$, or $t - 1$ otherwise.




Figure 1: Mapping of segments $\{[Y_j, Y_{j+1})\}_{j=0}^R$ to the points of the array X

4.4. Dealing with Floating Point Rounding Errors

Unfortunately the use of floating point arithmetic introduces a number of flaws in the procedure for the determination of H and R proposed in section 4.1. Some examples where computations fail are illustrated below.

- Given the partition $X = \{-10^9, 0, 1\}$ in single precision, $X_2 - X_0$ gets rounded to 10^9 and it is effectively indistinguishable from $X_1 - X_0$. Given function (4), regardless on the choice of H , it is impossible for property (1d) to hold.
- Given the array $X = \{0, 1.42 \cdot 10^{-45}, 1\}$ in single precision, equation (9a) overflows and yields $R = +\infty$.

Before continuing the discussion, it is useful to introduce some well known facts and notations about floating point numbers (for a more comprehensive discussion on this topic see [5]).

4.4.1. Facts and Notations about Floating Point Numbers

Floating point numbers are a finite subset of rational numbers and in a floating point numerical system any real numbers is approximated with its closest floating point number.

The result of an algebraic operations might not be a floating point numbers even if its operands are floating point numbers and it is therefore affected by a rounding error.

Amongst real positive numbers, only those bounded in a certain interval can be approximated with floating point numbers. Numbers too small or too big respectively underflow to zero or overflow to the abstract concept of $+\infty$.

Let:

- $m(z)$ be the floating point approximation of a real number z
- $\phi(x)$ be the smallest floating point number greater than the floating point number x
- ϵ be the *round-off error* associated with a given floating point representation, which is $\epsilon = 2^{-24}$ in single precision and $\epsilon = 2^{-53}$ in double precision

The relative rounding error made when a real number x is approximated by the floating point number $m(x)$ is bounded by

$$1 - \epsilon \leq \frac{m(x)}{x} \leq 1 + \epsilon \quad (10)$$

4.4.2. Feasibility Conditions

The algebraic operations involved in the computation of (4) and (9) are in general affected by rounding errors, therefore it is not guaranteed that the value of H they produce satisfies numerically condition (5). Taking into account rounding errors, conditions (5) become

$$\lfloor m(H m(X_{i+1} - X_0)) \rfloor > \lfloor m(H m(X_i - X_0)) \rfloor, \quad i = 0 \dots N - 1 \quad (11)$$

A necessary condition for (11) to hold is that the results of all arithmetic subtractions involved must be numerically distinguishable, i.e. strictly increasing:

$$m(X_{i+1} - X_0) > m(X_i - X_0), \quad i = 0 \dots N - 1 \quad (12)$$

Using (10) and taking repeatedly lower bound of the left hand side and a upper bound of the right hand side, conditions (12) become

$$\begin{aligned} (X_{i+1} - X_0)(1 - \epsilon) &> (X_i - X_0)(1 + \epsilon), & i = 0 \dots N - 1 \\ X_{i+1} - X_i &> (X_{i+1} + X_i - 2X_0) \epsilon, & i = 0 \dots N - 1 \\ \min_i \{X_{i+1} - X_i\}_{i=0}^{N-1} &> (2X_N - 2X_0) \epsilon \\ \frac{\min_i \{X_{i+1} - X_i\}_{i=0}^{N-1}}{X_N - X_0} &> 2\epsilon \end{aligned} \quad (13)$$

Assuming conditions (12) hold, conditions (11) can be rewritten as a function of the rounded interval lengths $\{D_i = m(X_i - X_0)\}_{i=0}^N$

$$\lfloor m(H D_{i+1}) \rfloor > \lfloor m(H D_i) \rfloor, \quad i = 0 \dots N - 1 \quad (14)$$

the truncation operation can be resolved writing a more restrictive set of inequalities,

$$m(H D_{i+1}) > m(H D_i) + 1, \quad i = 0 \dots N - 1 \quad (15)$$

and the theoretical lower bound (7) can be approximated numerically as

$$H > \bar{H} = m \left(\frac{1}{\min_i \{m(D_{i+1} - D_i)\}_{i=0}^{N-1}} \right) \quad (16)$$

The size of the array K is $R = \lfloor H X_N \rfloor$. In order to avoid numerical overflows, it must be $R < 2^Q$, where Q is the number of bits used for the index returned by function (4). For the function (4) to be efficiently vectorizable, it is desirable that $Q = 32$ or $Q \in \{32, 64\}$ depending if the array X is in single precision or in double precision. Note that Q is not necessarily the same as the number of bits used to represent elements of the array K , which is discussed in section 4.5. This requires that

$$m(H D_N) < 2^Q \quad (17)$$

Ignoring rounding errors, which might force H to be slightly larger than the theoretical lower bound (7), this limitation can be approximately expressed in terms of the layout of the original array X as

$$\frac{X_N - X_0}{\min_i \{X_{i+1} - X_i\}_{i=0}^{N-1}} < 2^Q \quad (18)$$

and the approximate feasibility conditions (13) and (18) can be combined in one single expression

$$\frac{\min_i \{X_{i+1} - X_i\}_{i=0}^{N-1}}{X_N - X_0} > \max \{2^{-Q}, 2\epsilon\} = \begin{cases} 2^{-23}, & \text{in single precision, with } Q = 32 \\ 2^{-32}, & \text{in double precision, with } Q = 32 \\ 2^{-51}, & \text{in double precision, with } Q = 64 \end{cases} \quad (19)$$

which approximately defines the family of arrays X where the method is applicable.

Note that conditions (13) and (18) are purely theoretical and cannot be verified exactly, because of rounding errors. It is however easy and inexpensive to verify directly the original conditions (12) and, given a value of H , (17).

4.4.3. Computation of a Feasible H

Both the value for H computed in (16) and the evaluation of function (4) are affected by rounding error, therefore there is no guarantee that choosing $H = \bar{H}$ conditions (14) are satisfied. Should they not be, a larger value of H is needed.

As discussed in section 4.1, it is desirable for H to be as small as possible. That poses the difficult question of *how much larger should H be?* There is no obvious answer and in the sequel a feasible value of H is computed numerically, by increasing H in small amounts in a trial and error iterative process. In brief, after a value for H is chosen, the following two steps happen in a loop: condition (15) is tested for all i and, if it does not hold, H is progressively increased according with some growth strategy.

A possible growth strategy for H is described in algorithm 9, which increments H adding terms of exponentially increasing size, until a feasible value is found.

Algorithm 9 Computation of H and R (pseudo-code)

```

function COMPUTEHR(input:  $\{X_i\}_{i=0}^N$  output:  $H, R$ )
     $H \leftarrow \phi(\bar{H})$  ▷ Initialize  $H$  strictly larger than the approximate lower bound (16)
     $D_N \leftarrow X_N - X_0$ 
    if  $\lfloor H D_N \geq 2^b \rfloor$  then ▷ Check for overflow verifying condition (17)
        ERROR: overflow, problem unfeasible
    end if
     $P \leftarrow \phi(H) - H$  ▷ Define a growth term  $P$ 
    for  $i = 1 \dots N$  do
         $D_{i-1} \leftarrow X_{i-1} - X_0$ 
         $D_i \leftarrow X_i - X_0$ 
        if  $D_{i-1} = D_i$  then ▷ Verify that the sequence  $D_i$  is strictly increasing
            ERROR:  $D_i$  are not strictly increasing, problem unfeasible
        end if
        while  $\lfloor H D_{i-1} \rfloor = \lfloor H D_i \rfloor$  do ▷ Check if condition (14) is satisfied
             $H \leftarrow H + P$  ▷ Increase  $H$ 
            if  $\lfloor H D_N \geq 2^b \rfloor$  then ▷ Check for overflow verifying condition (17)
                ERROR: overflow, problem unfeasible
            end if
             $P \leftarrow 2P$  ▷ Double the growth term  $P$ 
        end while
    end for
     $R \leftarrow \lfloor H X_N \rfloor$  ▷ Compute  $R$ 
end function

```

A minor modification of algorithm 9 would easily allow to also track the last unfeasible value for H , thus defining an interval $[H_{unfeasible}, H_{feasible}]$ which is known to be unfeasible at its left extreme and feasible at its right extreme. Then instead of simply taking $H = H_{feasible}$, the solution could be refined further by bisecting this interval. This has not been implemented in this paper, as in all scenarios tested in section 5.4, H is extremely rarely increased, and, when it is, the increment is negligible in relative terms.

4.5. Memory Cost

Elements of the array K must have a size B in bytes sufficiently large to store the largest index of the array X , i.e. N . To allow efficient memory manipulation, it should be $B \in \{1, 2, 4, 8\}$, i.e.

$$B = \min \left\{ b \in \{1, 2, 4, 8\} : 2^{8b} \geq N \right\}$$

Algorithm 8 requires the allocation of $R \cdot B$ bytes. The size R of the array K , apart from numerical related considerations, is approximately (9a), therefore, the total memory cost in bytes can be estimated with a quite high degree of accuracy as

$$MemoryCost \approx \left\lceil \frac{X_N - X_0}{\min_i \{X_{i+1} - X_i\}_{i=1}^N} \right\rceil B$$

4.6. Initial Setup Cost

The initial setup is divided in two parts, the computation of H performed in algorithm (9) and the initialization of the array K carried out in algorithm (7).

The latter has computational complexity $O(N + \alpha R)$, as it requires evaluation of function (4) for all elements of the array X and an assignment for all elements of arrays K . The constant of proportionality α reflects the fact that the two operations do not have the same cost and it is $\alpha \ll 1$.

It is more difficult to estimate precisely the computational complexity associated with the computation of H , as it depends on the number of iterations of the inner loop in algorithm (9), which increases H if needed. In theory it is of order $O(N(1 + T))$, where T is the average number of iterations in the inner loop. In practice, experimental results in section (5.4) show that $T \approx 0$, so the overall complexity is approximately $O(N)$.

The last component of the setup cost is the memory allocation of the array K , which is dependent on the programming language, the operating system and the memory allocation strategy used.

In section (5.4) some test results for the total setup cost, inclusive of memory allocation using the default *gcc* heap allocator, are reported.

4.7. Mitigating Limitations and Reducing Memory Consumption

It is possible to reduce the memory consumed by the index and mitigate limitations by relaxing the requirements of properties (1d) as follows

$$f(X_{i+2}) > f(X_i), \quad i = 0 \dots N - 2 \quad (20)$$

The definition of the index K needs to be generalized as:

$$K_j = \max\{i : f(X_i) \leq j\} \quad (21)$$

From definition (21), it follows the properties

$$a = f(X_i) < f(X_{i+1}) = b \implies K_a = i \quad (22a)$$

$$a = f(X_i) = f(X_{i+1}) = b \implies K_a = K_b = i + 1 \quad (22b)$$

Note that the modified definition of the array K does not require any change to algorithm (9).

Given a real number $z \in [X_i, X_{i+1})$, the sought index i can be resolved as

$$i = K_j - I(z < X_t) - I(z < X_{t-1}) \quad (23)$$

where $j = f(z)$ and $I(w)$ is the indicator function

$$I(w) = \begin{cases} 1, & \text{if } w \text{ is true} \\ 0, & \text{otherwise} \end{cases}$$

The proof of this statement is more complicated than in the case discusses in section 4. Property (1c) of function f guarantees that

$$a = f(X_i) \leq j = f(z) \leq f(X_{i+1}) = b,$$

which implies

$$K_a \leq t = K_j \leq K_b, \quad (24)$$

There are three possible scenarios:

1. if $f(X_i) < f(X_{i+1}) < f(X_{i+2})$, from property (22a) it follows that $K_a = i$, $K_b = i + 1$, and (24) implies $t \in \{i, i + 1\}$
2. if $f(X_i) < f(X_{i+1}) \leq f(X_{i+2})$, from properties (22) it follows that $K_a = i$, $K_b = i + 2$, and (24) implies $t \in \{i, i + 1, i + 2\}$
3. if $f(X_i) \leq f(X_{i+1}) < f(X_{i+2})$, from properties (22) it follows that $K_a = i + 1$, $K_b = i + 2$, and (24) implies $t \in \{i + 1, i + 2\}$

In all cases, t can only have one of the values in the set $\{i, i + 1, i + 2\}$:

- if $t = i$, equation (23) yields: $i - I(z < X_i) - I(z < X_{i-1}) = i - 0 - 0 = i$
- if $t = i + 1$, equation (23) yields: $i + 1 - I(z < X_{i+1}) - I(z < X_i) = i + 1 - 1 - 0 = i$
- if $t = i + 2$, equation (23) yields: $i + 2 - I(z < X_{i+2}) - I(z < X_{i+1}) = i + 1 - 1 - 1 = i$

which completes the proof.

Ignoring rounding errors, condition (6) becomes

$$H(X_{i+2} - X_0) > 1 + H(X_i - X_0), \quad i = 0 \dots N - 1 \quad (25)$$

yielding a smaller lower bound on H

$$H > \frac{1}{\min_i \{X_{i+2} - X_i\}_{i=0}^{N-2}} \quad (26)$$

which results in a smaller value of R and therefore requires a smaller storage space for the auxiliary array K . Note that algorithm 9 needs to be modified accordingly.

Limitations are also mitigated as conditions (19) become

$$\frac{\min_i \{X_{i+2} - X_i\}_{i=0}^{N-2}}{X_N - X_0} > \max \{2^{-Q}, 2\epsilon\} = \begin{cases} 2^{-23}, & \text{in single precision, with } Q = 32 \\ 2^{-32}, & \text{in double precision, with } Q = 32 \\ 2^{-51}, & \text{in double precision, with } Q = 64 \end{cases} \quad (27)$$

which are clearly less restrictive.

Summarizing, given the function f and the array $\{K_j\}_{j=0}^R$, the search procedure is described in algorithm 10. The algorithm is very similar to the one previously described, with the difference

Algorithm 10 Direct Search Gap2 (scalar implementation)

```

function DIRECTSEARCHGAP2(input:  $z, \{X_i\}_{i=0}^N, \{K_j\}_{j=0}^R, H$ , output:  $i$ )
     $j \leftarrow \lfloor H(z - X_0) \rfloor$ 
     $i \leftarrow K_j$ 
    if  $z < X_t$  then
         $i \leftarrow i - 1$                                  $\triangleright$  conditional assignment
    end if
    if  $z < X_{t-1}$  then
         $i \leftarrow i - 1$                                  $\triangleright$  conditional assignment
    end if
end function

```

that the number z needs to be compared with both X_t and X_{t-1} , with consequent performance degradation, and the index i could be accordingly decreased by 1 or 2. The second conditional assignment could be made contingent on the first condition $z < X_t$, but in this case the first condition would no longer just trigger a conditional assignment and become a genuine code flow branch, hence it is preferable to check the two conditions independently. The downside is that, if $t = 0$, this causes access to the array element X_{-1} , therefore the array X need to be padded to the left with one extra element containing the value X_0 . This does not affect the correctness of the algorithm because the condition $z < X_0$ is always false.

Note that this same idea could be pushed even further, generalizing condition (1d) to

$$f(X_{i+q}) > f(X_i), \quad i = 0 \dots N - q \quad (28)$$

which requires comparisons of z with $X_t, X_{t-1}, \dots, X_{t-q}$.

This would allow to construct an optimal solution deriving the minimal number of q which achieves a given memory budget.

4.8. Cache Efficient Implementation

Noting that in algorithm 11 the following two steps always occur in close sequence:

1. the index $t = K_j$ is read from the array K

2. the element X_t is read from the array X

suggests to reorganize the data in memory so that K_j and X_t are stored contiguously in memory and part of the same cache line ⁶. In this way, whenever K_j is fetched, X_t is loaded in cache at the same time and it is already available when it is needed for the execution of the next instruction, thus reducing the number of *cache misses*.

This is easily achieved modifying the data type stored in the array K to be the pair of values (K_j, X_t) , appropriately padded so that their joint memory storage requirement is either 8 or 16 bytes to guarantee good cache line alignment. For instance working in single precision with 32-bit indices, the pair requires exactly 8 bytes, while working in double precision with 32-bit indices, the pair requires 12 bytes and an extra 4 bytes padding needs to be appropriately inserted. Obviously this algorithm has higher memory requirements than algorithm 8, as elements of the array X are stored multiple times, plus there may be some wasted space for padding.

The pseudo-code is illustrated in algorithm 11.

Algorithm 11 Direct Search Cache Friendly (scalar implementation)

```

function DIRECTSEARCHWITHCACHE(input:  $z, \{X_i\}_{i=0}^N, \{K_j\}_{j=0}^R, H$  output:  $i$ )
     $j \leftarrow \lfloor H(z - X_0) \rfloor$ 
     $(i, x) \leftarrow K_j$  ▷  $K_j$  contains the pair of values  $(i, X_i)$ 
    if  $z < x$  then
         $i \leftarrow i - 1$  ▷ conditional assignment
    end if
end function

```

5. Test Results

5.1. Data Types Used

All tests results presented in this section are performed for arrays X in *single* and *double* precision. For simplicity of implementation, the indices used in algorithm 8, 10 and 11 are always 32-bits *unsigned integers*, regardless of the size of array K . Elements of the array K are always 32-bits *unsigned integers*, regardless of the size of the array X .

5.2. Array X Layout

All tests results presented in this section are based on arrays X of size $2^P - 1$, for $P = \{4, 8, 12, 16, 20\}$. The choice of this particular set of array sizes, near to powers of 2, is simply aimed at optimizing memory usage for algorithms 3 and 6 and has no impact on performance measurements for algorithms 8, 10 and 11.

Regardless of the size, all arrays X are composed of sub intervals of random length $(X_{i+1} - X_i)$ sampled for each i from a uniform distribution in $[1, 5]$.

The motivation associated with the choice of this particular layout is merely to assure that the feasibility condition (19) for algorithm 8 is satisfied in probabilistic terms even for the largest size of the array X tested. In fact, the expected length of an interval $\mathbb{E}[X_{i+1} - X_i]$ is 3 and, when

⁶a typical cache line is 64 bytes long and starts at memory addresses which are multiple of 64

the size of the array X is $N \approx 2^{20}$, condition (19) holds with good margin both in single and double precision

$$\frac{\min_i \{X_i - X_{i-1}\}_{i=1}^N}{X_N - X_0} \geq \frac{1}{X_N - X_0} \approx \frac{1}{\mathbb{E}[X_N - X_0]} = \frac{1}{3 \cdot 2^{20}} = 1.5 \cdot 2^{-21} > 2^{-23}$$

Furthermore, since the elements of the array Z are uniformly distributed in $[X_0, X_N)$, the choice of a layout X with intervals of comparable length, assures that there is no major disproportion in the number of times each element of the array X is picked, thus avoiding test results to be affected by cache memory related effects.

Alternative layouts could be tested, and, with respect to algorithms 8, 10 and 11, they might lead to partitions either feasible or unfeasible, but this would make no difference for the performance measurements carried out, except perhaps for considerations associated with cache memory.

5.3. Performance Tests

5.3.1. Performance Test Details

Tables 1-5 compare the throughput of various algorithms for arrays X of increasing size, as described in section 5.2. Results obtained using the Intel Math Kernel Library (MKL), which uses a proprietary algorithm, are also included ⁷.

For most algorithms there are 3 possible implementations for each floating point representation, where the set of instructions used is either scalar or vectorial with vectors of 128 bits (Intel SSE instruction set) or 256 bits (Intel AVX instruction set). Exceptions are algorithms 1, where the vectorial implementation consists in simply calling in a loop the scalar implementation with *function inlining* enabled, and *MKL*, where the scalar implementation consists in calling the vectorial implementation with vectors of size 1 and there is no way to force use of only SSE or AVX instructions. In both cases the same identical algorithm is used to produce the numbers reported in the SSE and in the AVX section.

For each floating point representation and set of instructions the results are expressed as the number of indices resolved per unit of time. The setup-cost, for those algorithms requiring it, is excluded from the reported results.

The array Z is stored aligned on a 32 bytes boundary, for efficient vectorial memory access. Its elements have no particular pattern and are randomly uniformly distributed in $[X_0, X_N)$.

5.3.2. Performance Test Observations

MKL. MKL performance is generally superior to classic binary search, but significantly inferior to any of the enhanced versions of binary search discussed. The scalar case in particular, for small arrays, is even slower than classic binary search.

Enhanced Versions of Binary Search. The three enhanced versions of binary search (algorithms 3, 5 and 6) for arrays of small to medium size exhibit comparable performance in the vectorial case and only minor differences in the scalar case, with algorithm 3 being faster. For large arrays, as expected, the cache friendly algorithm 6 beat the others.

⁷The version of MKL used is 2017. The functions used are *dfsSearchCells1D* in single precision and *dfdSearchCells1D* in double precision. The partition X is initialized with the *QUASLUNIFORM* hint, which yields experimental results superior to the alternatives for the particular layouts of the array X used in this test

Direct Search. All versions of direct search (algorithms 8, 10, 11) yield a throughput increase of at least one order of magnitude with respect to binary search algorithms. This is no surprise given the difference in computational complexity. As expected, algorithm 10 has slightly inferior performance than algorithm 8, due to the increased computation cost. Algorithm 11 performance is superior to all others for large vectors, where an efficient utilization of cache memory becomes critical.

General Observations. All algorithms show a significant speed-up switching from the scalar to the SSE vectorial implementation, while switching from SSE to AVX leads to less consistent improvements, which vary with the algorithm, the precision and the array size. Note that the throughput increase factor achievable via vectorization is expected to be smaller than d . This is because, using Intel SIMD instructions, as explained in previous sections, perfect vectorization is not achievable.

5.4. Direct Search Setup Cost Test

5.4.1. Direct Search Setup Cost Test Details

As discussed in section 4.6, the setup cost is uncertain because of the unpredictable number of times H is increased in the inner loop of 9. Therefore all test results in this section are described with statistical properties over populations of randomly generated arrays X .

Table 6 presents the number of times H is increased in the inner loop of algorithm 9. A sample population of 10000 randomly generated arrays X is used.

Table 7 presents the setup cost in nanoseconds normalized by the size of the array X . A sample population of 1000 randomly generated arrays X is used.

Multiplying the test results in table 7 for the array size yields the average setup time in nanoseconds. This allow for a direct comparison with the results presented in tables 1-5, which are expressed in million of searches per seconds, making it possible to express the setup costs in terms of equivalent number of searches.

5.4.2. Direct Search Setup Cost Test Observations

Table 6 shows that the number of times H is increased in the inner loop of algorithm 9 is independent on N . At most one iteration is carried out, regardless of the size of the array X . This implies that the initial guess (16) is acceptable in almost all cases and, when an increase is necessary, that is negligible in relative terms. Hence the approximate limiting conditions (19) which is derived from (16) have good general accuracy.

Table 7 shows that the setup cost normalized by the size of the array X , initially decreases as the size of the array X grows, but then it increases until it stabilize on an asymptotic level.

The initial decrease can be explained as the effect of the fixed overhead associated with the algorithm, due for instance to the allocation of memory for the array K , which is material for a small array, but gets amortized over larger arrays.

The subsequent increase can be explained by the larger amount of total memory necessary to store the arrays X and K , which results in progressively inefficient use of the cache memory.

From these experimental results it is possible to conclude that, a part from memory I/O effects, the setup cost per element is independent on N , i.e. the setup has complexity $O(N)$.

5.5. Reproducibility

All test results presented in section 5 have been produced using the C++ source code publicly available from [github](https://github.com/fabiocannizzo/fastbinarysearch)⁸, compiled for Linux 64-bits platform with gcc 6.3.1 on a machine with AVX-2 capability⁹.

6. Conclusion

Possible improvements to the classic binary search algorithm have been described. Although having the same asymptotic complexity, these are vectorizable and generally faster.

Next a new algorithm with superior asymptotic complexity has been presented. Test results using streaming SIMD extensions demonstrate that with arrays X of various length randomly populated the proposed algorithm is up to 43 times faster than the classical binary search in single precision and up to 39 times faster in double precision.

A cache friendly version of the algorithm is also proposed, which in most cases leads to even superior performance.

There are situations where the algorithm is not applicable. These are extensively discussed and can be inexpensively identified in the preliminary analysis phase, thus allowing fall-back to some of the enhanced versions of the binary search algorithm. Possible variations of the algorithm, which mitigate such limitations sacrificing some performance or reduce cache misses at the cost of requiring more memory, have also been proposed.

References

7. References

- [1] 2011, Pulver, Binary Search Revisited, URL: <http://eigenjoy.com/2011/09/09/binary-search-revisited/>
- [2] 2015, P. Khuong, P. Morin, Array Layout for Comparison Based Searching, ARXIV
- [3] 1997, D. Knuth, Sorting and Searching, volume 3 of The Art of Computer Programming. Addison-Wesley, second edition.
- [4] 2007, W. Press, S. Teukolsky, W. Vetterling, B. Flannery, Numerical Recipes, The Art of Scientific Computing, Sections 3.1, Third Edition, Cambridge University Press
- [5] 1991, D. Goldberg, What Every Computer Scientist Should Know About Floating-Point Arithmetic, Computing Surveys, Association for Computing Machinery
- [6] 1982, L. Johnson, R. Riess, Numerical Analysis, Second Edition, Addison Wesley
- [7] 2008-754 IEEE Standard for Floating-Point Arithmetic, The Institute of Electrical and Electronics Engineers, 2008
- [8] 2012, P. Khuong, Binary search is a pathological case for caches, Some Lisp, URL: <http://www.pvk.ca/Blog/2012/07/30/binary-search-is-a-pathological-case-for-caches/>.
- [9] 1946, J. Mauchly, Theory And Techniques for Design of Electronic Digital Computers, Lectures given at the Moore School of Electrical Engineering
- [10] 1960, D. H. Lehmer, Proceedings of Symposia in Applied Mathematics 10
- [11] 1962, H. Bottenbruch, JACM 9
- [12] 1953, J. Kiefer, Sequential minimax search for a maximum, Proc. American Mathematical Society 4, 502506.
- [13] 1960, D. Ferguson, Fibonacci searching, Communications of the ACM, vol. 3, is. 12, p. 648
- [14] 1976, J. Bentley, A. Yao, An almost optimal algorithm for unbounded searching, Information Processing Letters. 5 (3): 8287
- [15] 1590, M. Eytzinger, Thesaurus principum hac aetate in Europa viventium (Cologne) Aitsingero, Aitsingerum, Eytzingerum.

⁸<https://github.com/fabiocannizzo/fastbinarysearch.git>

⁹CPU model: Intel(R) Xeon(R) CPU E5-2666 v3 @ 2.90GHz

	Single Precision			Double Precision		
	Scalar d=1	SSE-4 d=4	AVX-2 d=8	Scalar d=1	SSE-4 d=2	AVX-2 d=4
Algorithm 1	48.30	48.96	51.28	47.67	50.86	49.47
Algorithm 3	177.18	375.26	396.64	197.96	263.19	214.13
Algorithm 5	180.45	369.56	396.06	182.14	255.83	213.23
Algorithm 6	161.30	374.44	391.60	160.98	254.03	207.07
MKL 2017	34.55	167.08	167.55	31.70	152.89	151.93
Algorithm 8	331.11	726.62	728.05	330.27	492.58	553.76
Algorithm 10	207.78	576.57	607.32	207.08	410.46	453.70
Algorithm 11	328.13	732.89	754.83	312.39	570.22	469.42

Table 1: Throughput in million of searches per second with vector X of size 15

	Single Precision			Double Precision		
	Scalar d=1	SSE-4 d=4	AVX-2 d=8	Scalar d=1	SSE-4 d=2	AVX-2 d=4
Algorithm 1	20.55	20.40	20.50	20.45	19.94	20.46
Algorithm 3	80.49	122.18	148.41	81.29	89.98	81.36
Algorithm 5	73.47	122.87	148.01	74.05	90.36	81.15
Algorithm 6	65.63	122.45	146.56	65.25	88.94	79.97
MKL 2017	24.95	83.44	83.50	22.43	84.18	84.21
Algorithm 8	331.26	722.33	728.57	330.16	491.94	554.18
Algorithm 10	207.68	575.63	605.53	206.71	409.92	451.44
Algorithm 11	328.17	730.32	733.19	312.16	569.70	470.62

Table 2: Throughput in million of searches per second with vector X of size 255

	Single Precision			Double Precision		
	Scalar d=1	SSE-4 d=4	AVX-2 d=8	Scalar d=1	SSE-4 d=2	AVX-2 d=4
Algorithm 1	14.08	14.15	14.16	13.95	13.87	14.10
Algorithm 3	45.30	70.75	91.33	44.45	47.07	48.48
Algorithm 5	42.98	70.96	91.14	41.82	47.39	48.42
Algorithm 6	38.51	70.34	90.00	37.95	46.59	48.24
MKL 2017	21.57	53.52	53.54	19.23	51.65	51.65
Algorithm 8	309.78	642.61	618.90	302.05	427.90	487.01
Algorithm 10	195.31	544.69	548.09	186.98	352.28	396.78
Algorithm 11	302.19	705.73	639.64	259.23	476.40	387.07

Table 3: Throughput in million of searches per second with vector X of size 4095

	Single Precision			Double Precision		
	Scalar d=1	SSE-4 d=4	AVX-2 d=8	Scalar d=1	SSE-4 d=2	AVX-2 d=4
Algorithm 1	9.05	9.11	9.10	7.13	7.13	7.20
Algorithm 3	20.89	28.86	38.04	12.81	11.92	15.47
Algorithm 5	18.45	26.39	35.04	12.28	11.84	15.44
Algorithm 6	22.31	44.63	50.08	17.55	25.73	22.52
MKL 2017	15.39	26.21	26.21	13.12	22.89	22.89
Algorithm 8	159.67	345.02	349.01	148.88	246.81	281.38
Algorithm 10	118.57	315.81	331.39	108.98	199.31	241.49
Algorithm 11	239.46	547.16	511.86	218.39	394.57	333.06

Table 4: Throughput in million of searches per second with vector X of size 65535

	Single Precision			Double Precision		
	Scalar d=1	SSE-4 d=4	AVX-2 d=8	Scalar d=1	SSE-4 d=2	AVX-2 d=4
Algorithm 1	4.59	4.62	4.62	3.70	3.70	3.73
Algorithm 3	5.49	10.68	17.11	4.84	5.23	7.15
Algorithm 5	6.17	11.14	17.44	5.04	5.40	7.33
Algorithm 6	9.44	18.79	23.76	7.94	9.98	10.95
MKL 2017	7.83	9.99	9.98	6.78	8.74	8.74
Algorithm 8	62.46	61.68	59.37	53.02	52.68	50.69
Algorithm 10	67.08	75.00	72.22	57.75	60.51	58.73
Algorithm 11	105.96	96.31	92.64	82.17	82.43	76.13

Table 5: Throughput in million of searches per second with vector X of size 1048575

array size	Single Precision				Double Precision			
	mean	min	max	stdev	mean	min	max	stdev
15	0.0061	0.0000	1.0000	0.0779	0.0048	0.0000	1.0000	0.0691
255	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	1.0000	0.0100
4095	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
65535	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1048575	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 6: Number of H updates

array size	Single Precision				Double Precision			
	mean	min	max	stdev	mean	min	max	stdev
15	10.24	8.96	11.59	0.47	11.81	10.11	13.31	0.49
255	8.27	8.00	8.66	0.12	8.40	8.06	8.65	0.13
4095	13.80	13.54	13.98	0.09	13.51	13.30	13.70	0.07
65535	17.63	17.61	17.65	0.01	17.49	17.46	17.51	0.01
1048575	17.67	17.66	17.69	0.01	17.61	17.58	17.66	0.02

Table 7: Statistical setup cost for algorithm 8 in nano seconds normalized by the array size